Contents lists available at ScienceDirect

Reliability Engineering and System Safety

Quantifying reliability uncertainty from catastrophic and margin defects: A proof of concept

Christine M. Anderson-Cook [a,1], Stephen Crowder [b,2], Aparna V. Huzurbazar [a,1], John Lorio [b,*,2], James Ringland [c,2], Alyson G. Wilson [d]

[a] Statistical Sciences Group, Los Alamos National Laboratory, P.O. Box 1663, MS F600 Los Alamos, NM 87545, USA
[b] Sandia National Laboratories, Albuquerque, NM 87185, USA
[c] Sandia National Laboratories, Livermore, CA 94550, USA
[d] Department of Statistics, Iowa State University, Ames, IA 50011, USA

ABSTRACT

We aim to analyze the effects of component level reliability data, including both catastrophic failures and margin failures, on system level reliability. While much work has been done to analyze margins and uncertainties at the component level, a gap exists in relating this component level analysis to the system level. We apply methodologies for aggregating uncertainty from component level data to quantify overall system uncertainty. We explore three approaches towards this goal, the classical Method of Moments (MOM), Bayesian, and Bootstrap methods. These three approaches are used to quantify the uncertainty in reliability for a system of mixed series and parallel components for which both pass/fail and continuous margin data are available. This paper provides proof of concept that uncertainty quantification methods can be constructed and applied to system reliability problems. In addition, application of these methods demonstrates that the results from the three fundamentally different approaches can be quite comparable.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

Quantification of Margins and Uncertainties (QMU) analysis is an approach towards computing and expressing component margin relative to a requirement or relative to a threshold with a connecting component. A QMU analysis may be as simple as a histogram of a component's performance output shown relative to a performance requirement, with the margin expressed in standard deviations from the mean. In more complex applications, such as components that have an age trend in their performance data, the QMU analysis could involve a linear regression against component age, along with uncertainty bands around the regression line. In such QMU analysis we are generally interested in making an end of life prediction.

When executing a QMU analysis we can express a component's capability in terms of K factors. K is defined as (margin/uncertainty), or more explicitly as

$K = [(\text{upper threshold} - \text{sample mean})/\text{sample standard deviation}]$

for an upper threshold;

or

$K = [(\text{sample mean} - \text{lower threshold})/\text{sample standard deviation}]$

for a lower threshold.

K provides a common measure of margin. Given a distribution function fitted to the data, K also correlates to a fraction defective (that is, fraction that fails to meet the requirement). If one has data normally distributed (a big IF), K factors of 3 or greater are considered adequate. As we often have data that is not normally distributed, we examine an expected fraction defective instead in evaluating the adequacy of the margin.

A QMU analysis is based upon critical performance variables using data taken during lot acceptance, surveillance, flight testing, and other sources. A typical QMU analysis spends most of the initial resources in sorting through this data teasing out significant factors that may influence the data, such as changes to testers during the elapsed time the data was taken, trends during production, test conditions, and lot-to-lot variation.

Computing the margin (defined as data mean—requirement) requires some definition of the requirement. To date this has generally meant using a published requirement in an interface definition or product acceptance specification. However, we have lately advanced towards calculating joint probabilities between the performance distributions and threshold distribution of two connecting components in order to correctly understand the true margin between the components.

To date, the Nuclear Weapons complex has conducted many component level QMU analyses [1–5] but has not yet developed ways to integrate these component level analyses into a framework for making system level inferences. What impact does low margin of a component have on the overall system reliability and uncertainty? Being able to answer such a question allows us to leverage QMU assessments to improve system level decisions, as well as assess the impact of low component margins on overall system level reliability and uncertainty. We can do this by integrating margin insufficiency failure modes with quality defect failure modes in a common model. By computing the overall system uncertainty we get a measure of "confidence" in the system that we can use as a decision tool in evaluating different system projects and in identifying where additional testing resources should be allocated.

Nuclear weapons are designed with multiple objectives that include safety, security, and reliability. We limit our focus to reliability, defined as the probability of success of the weapon performing its intended function at the intended time given the required temperature range, shock and vibration exposures, altitude and speed of the release envelope and over the designed lifetime of the weapon. The goal of the Department of Energy (DOE)/National Nuclear Security Administration (NNSA) weapon reliability assessment process is to provide a quantitative metric for this assessment. More details about this process are given in Wright and Bierbaum [6].

We typically use both system level and component level test data in computing the system reliability, using a series/parallel model of the critical events to compute the overall probability of success. Component level failure modes (and data) lie behind each of the critical events. Thus several component level failure modes are combined using series/parallel models to compute the probability of an event occurring successfully. Component level failure modes are computed using failure data; thus the reliability of an event is some series/parallel combination of component reliabilities. In equation form, for two components joined in series for a particular event, we have

$$R_{event} = R_{component\ 1} * R_{component\ 2}$$
$$= (1 - Prob(component\ 1\ failure))(1 - Prob(component\ 2\ failure))$$

This paper demonstrates a proof of concept of three methods for estimating system level reliability and uncertainty using component-level data. The data used in the analyses presented here consist of both catastrophic failures (pass/fail data) and margin failures (continuous data). The analyses are applied to a relatively simple system model, consisting of both mixed series and parallel components. Three diverse methods (Method of Moments, Bayesian, and Bootstrap) are used and the results from each are compared to the NNSA system reliability point estimate approach [7]. The example system model captures some key features of the top-level models used by Sandia National Laboratories and Los Alamos National Laboratory to assess weapon reliability. While this example uses a two level model (events and component failure modes), all three approaches can be applied on more complex systems (though with varying degrees of difficulty) and are not restricted to two levels.

The NNSA approach, described in Section 2, is currently the standard method used for reporting point estimates for system reliability. Historically, point reliability estimates have sufficed to support military operational planning. The NNSA Surveillance Program has historically focused on detecting quality defects early in the life of the stockpile rather than uncertainty measures associated with the known failure modes and existing data.

With a mature stockpile, limited production opportunities, and a now-extensive surveillance data base, understanding the residual uncertainty associated with known and measurable failure modes has grown in importance. A key feature of QMU reliability analyses is to focus not only on a point estimate of reliability, but also on the uncertainty associated with the estimate. Understanding uncertainty can contribute to subsequent decision-making. Hence the three methods presented here seek to complement the NNSA point reliability and to provide a unified mechanism for assessing system-level reliability (point estimates) and reliability uncertainty (interval estimates) associated with component-level catastrophic failures and margin failures. Each method can measure the contribution of each failure source to the overall system reliability uncertainty. Each method can also be applied to a variety of system structures including mixed series/parallel systems.

The first method that complements the NNSA point estimate by adding appropriate uncertainty intervals is a classical Method of Moments (MOM) approach. (See Kotz et al. [8] for a summary discussion of the Method of Moments and an extensive list of historical and more recent references.) This approach captures and aggregates sampling uncertainties at the component level. The method evaluates the mean and variance of the various component-level reliability estimators and then propagates these through the system level reliability equation. Simplifying assumptions are made as to the distributional form of the system reliability distribution to get an approximate value for the mean and variance of the system reliability estimator. We then construct a 90% confidence interval for the system reliability using an equivalent Binomial distribution.

The second method presented applies some of the modern computing-intensive approaches currently available. It is a Bayesian approach using Markov Chain Monte Carlo (MCMC) methods to develop the system level distribution for reliability. This approach selects a user-specified diffuse prior distribution for each component, which is then updated based upon the data to develop the posterior distribution. By this mechanism, both sampling and other knowledge uncertainties can be captured, although the specific example here emphasizes the former. The individual component level reliability estimates are then combined to obtain the system level reliability estimate with an associated uncertainty propagated from the component distributions. We then construct a 90% credible interval for system reliability using the computed system-level reliability distribution.

The third method also takes advantage of modern computing power and considers resampling through a Bootstrap approach. It re-samples the available data with replacement to develop the non-parametric component and system level probability distributions. The resulting system-level distribution is used to construct a 90% Bootstrap confidence interval for system reliability. This method employs the fewest assumptions of the three approaches.

All the three approaches begin with a small number of basic principles but then require a fair amount of mathematical machinery to execute. All three can be complex, depending on the complexity of the system under study. The MOM approach involves substantial analytic manipulation of the equations. The Bayesian approach is computationally intensive and relies on MCMC simulation, with tools becoming more readily available in software packages, such as R and WinBUGs. The Bootstrap approach is also computationally intensive, but can be programmed into readily available software such as Minitab or Matlab.